

AD-A260 384



2

A RAND NOTE

Exploratory Modeling and the Use of Simulation for Policy Analysis

Steven C. Banks



DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

RAND

93-03070



4120

The research described in this report was sponsored by the United States Army under Contract No. MDA903-91-C-0006.

Accession For	
NTIS (GAI)	<input checked="checked" type="checkbox"/>
DTIC (GAI)	<input type="checkbox"/>
Unreviewed	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Serial and/or	
Dist	Serial
A-1	

DTIC QUALITY INSPECTED 3

RAND is a nonprofit institution that seeks to improve public policy through research and analysis. Publications of RAND do not necessarily reflect the opinions or policies of the sponsors of RAND research.

A RAND NOTE

N-3093-A

**Exploratory Modeling and the Use of
Simulation for Policy Analysis**

Steven C. Bankes

**Prepared for the
United States Army**

RAND

Approved for public release; distribution unlimited

PREFACE

This Note documents insights into the use of computer modeling for difficult policy problems. These insights have been arrived at in the context of the project "Measuring the Operational Value of Intelligence and Electronic Warfare (OPVIEW)." The project, part of the Applied Technology Program, was undertaken by the Army's Arroyo Center. Its objective is to develop a methodological approach and supporting model as tools for analyzing and measuring the operational value of Intelligence and Electronic Warfare (IEW) in combat outcome terms. This Note expresses a general "vision" of how to use exploratory modeling to support policy studies. To see how this vision is being applied to analyzing the contribution of IEW to combat outcomes, see N-3101-A, *Methodological Considerations in Using Simulation to Assess the Combat Value of Intelligence and Electronic Warfare*.

THE ARROYO CENTER

The Arroyo Center is the U.S. Army's federally funded research and development center (FFRDC) for studies and analysis, operated by RAND. The Arroyo Center provides the Army with objective, independent analytic research on major policy and management concerns, emphasizing mid- to long-term problems. Its research is carried out in five programs: Policy and Strategy; Force Development and Employment; Readiness and Sustainability; Manpower, Training, and Performance; and Applied Technology.

Army Regulation 5-21 contains basic policy for the conduct of the Arroyo Center. The Army provides continuing guidance and oversight through the Arroyo Center Policy Committee, which is co-chaired by the Vice Chief of Staff and by the Assistant Secretary for Research, Development, and Acquisition. Arroyo Center work is performed under contract MDA903-91-C-0006.

The Arroyo Center is housed in RAND's Army Research Division. RAND is a private, nonprofit institution that conducts analytic research on a wide range of public policy matters affecting the nation's security and welfare.

Lynn E. Davis is Vice President for the Army Research Division and Director of the Arroyo Center. Those interested in further information concerning the Arroyo Center

should contact her office directly:

Lynn E. Davis

RAND

1700 Main Street

P.O. Box 2138

Santa Monica, CA 90406-2138

Telephone: (310) 393-0411

SUMMARY

The use of computer models for most policy analysis purposes has a fundamentally different character from what is classically considered modeling in engineering and the hard sciences. Models in the physical sciences are often used to make detailed predictions; most systems of concern to policymakers, however, cannot be predictively modeled. This Note describes how "exploratory modeling" provides a rationale for how computer models can be fruitfully employed in support of policy studies. It also suggests improvements that could be made both in the methodology and technology of computer modeling for policy.

The use of computer models for policy analysis has encountered a variety of difficulties, including:

- Unwieldy size
- Problems with verification and validation
- Inadequate sensitivity analysis
- User inability to understand model internals and outputs
- Systematic bias arising from considering only those phenomena for which good models can be constructed.

Potentially undesirable consequences of these difficulties include using models to rationalize institutional prejudices, poor models driving out careful thinking, and tending to emphasize those aspects of a problem that can best be simulated. The result can often be that computer models provide an illusion of analytic certainty for problems that are not that well understood or, in the worst cases, provide scientific costume for points of view that are self-serving. Many of these problems are the result of a confusion between predictive modeling and the exploratory use of models to support reasoning about complex problems.

Most models used for policy studies will have significant associated uncertainties that bar their use for prediction and make validation of model correctness impossible. When experimental validation is no longer possible, no model can be asserted to be a correct model of the target system, and no conclusions can be safely drawn without examining alternative models.

For problem areas in which prediction is not possible, computer models may nevertheless be useful as a means to assist reasoning. Even quite simple models are capable of exhibiting behaviors that surprise their creators. Consequently, unanticipated implications of our assumptions may be revealed by building and executing models. Even when a model is not validated, it can serve as an "inference engine," showing us where innocuous appearing assumptions lead to predicted behaviors at variance with initial expectations. By throwing light on obvious but treacherous assumptions, computer modeling can perform an important service in the search for understanding. Thus, the use of models can help us to discover novel insights, although establishing the relevance of any such insight to particular problems must be accomplished outside of the models.

In exploratory modeling, a model run cannot be considered a prediction of events in the world; it is rather a computational experiment that yields information about the model itself. Where prediction and experimental validation is not possible, no single "true" model for the system of interest can be agreed upon. However, in such situations we need not be restricted to the use of any single model in our reasoning. Although all plausible models cannot be examined, insight can be gained by exploration of a number of them.

The goal of exploration is a compelling argument illuminating the choice among policy options. In constructing such an argument, models must be built and used in service to an analytic strategy, and in a study's conclusions they are relevant only in the context of an argument that takes their limitations into account. The unaugmented outputs of a nonpredictive model never have a particular meaning for policy.

Because a thorough search among all plausibility models is not possible, analytic strategies will often focus on the elucidation of critical cases that support choosing one policy from a list of options even when the exact range of outcomes is unknown. Examples are the discovery of taxonomies of worst cases that allow choosing among options by risk aversion, the use of a fortiori arguments, and identifying cases where expenditures make a difference or where competing models may be distinguished. In this way, even where models do not predict, they can be used to discover facts true of all plausible models or to convince by example. However, their value is always relative to a logical context established outside the models. The argument is the central result, not the outputs of any computer program.

Exploratory modeling allows for the flexible allocation of resources (human as well as computational) to those aspects of the problem that are judged most critical in arriving at such an argument. Given a fixed analytic budget (in dollars, people, or time), analysis must provide the best insights possible based on what is known about the problem at hand.

The increasing availability of computer power makes aggressive pursuit of exploratory modeling a growing possibility. However, this will require changes in both the technological support for computer modeling and the methodology of model development. This Note suggests several innovations in the methodology of model construction that would facilitate better exploitation of exploratory modeling's potential. These innovations include:

- Model design driven by the question being asked rather than by details of the system being studied. Rather than building a single model which is used to answer a variety of questions (as would be possible for predictive, experimentally validated models), the model is designed to answer a specific question. Models made to serve study goals will result in better studies than studies driven by what can be modeled.
- The use of multiple models rather than a single "monolithic" model. Since no model generated will be a valid representation of the true state of the world, there is no need to produce a single model. Consequently, the complexities of the target system can be represented by an ensemble of models and by supporting analysis. The scope of any individual model can be designed to maximize its utility for answering a specific question (including the need for understandability and sensitivity analysis). The linear process of developing a conceptual model, implementing it as a program, and running cases is replaced by a much more iterative approach, in which preliminary modeling leads to insights that affect the design of subsequent models.
- Model development by a process of "selective resolution." This approach uses aggregate models for preliminary analysis that then guides the development of more detailed models. In particular, the aggregate models are used to discover what factors are critical to deciding a particular question, allowing only the most critical factors to be disaggregated in subsequent modeling.

The Note also suggests improvements to computer technology (in computer language, human interface, and complexity management tools as part of an integrated simulation environment) to support modeling. The methodology proposed here would in many cases replace a complex monolithic model with numerous simpler models. Relative model simplicity would be obtained by exporting complexity outside the model and into the surrounding computational environment and analytic context. One reason why this shift is

attractive is that when complexity is not buried inside a model, more powerful software tools can be provided to manage it. Aggressive exploitation of the approach championed here requires the development of such tools. These tools taken together amount to a computer environment to support exploratory modeling. Such an environment should allow smart people to navigate efficiently through the space of plausible models and model outcomes to construct lines of reasoning and to make themselves smarter. Proposed computer tools would meet the following three general needs:

- Support for iterative and adaptive modeling
- Assistance in managing the complexity of numerous models, cases, and relationships between them
- Means for portraying the results of exploratory modeling.

The use of exploratory modeling cannot make the uncertain certain, nor can it make complex problems simple. But it can motivate better use of computers in support of policy analysis, provide for a better allocation of resources in dealing with the real problems, and afford some protection against fooling ourselves.

ACKNOWLEDGMENTS

The ideas articulated here have developed in the context of vigorous research into simulation at RAND, and many of my colleagues have contributed in one way or another. The author appreciates the thoughtful review comments of Patrick Allen, Edison Cesar, Edward Hall, James Hodges, Iris Kameny, Eric Larson, Jed Marti, Norman Shapiro, Warren Walker, and Robert Worley.

CONTENTS

PREFACE	iii
SUMMARY	v
ACKNOWLEDGMENTS	ix
Section	
1. INTRODUCTION	1
2. THE PROBLEM: PRETENDING TO DO WHAT CAN'T BE DONE	3
The Shortcomings of Computer Models	3
Why Models Go Wrong	6
Trying to Predict the Unpredictable	8
3. EXPLORATORY MODELING	11
The Search for Insight	11
Sensitivity Analysis	13
Using Exploratory Modeling to Support Policy Analysis	14
4. IMPLICATIONS FOR BUILDING EXPLORATORY MODELS	17
Question-Driven Modeling versus Data-Driven Modeling	17
Not One Model, but Many	18
Selective Resolution	19
Computer Technology to Support Exploratory Modeling	20
5. SCIENCE, OPERATIONS RESEARCH, AND THE ROLE OF THE COMPUTER	26
BIBLIOGRAPHY	29

1. INTRODUCTION

The use of computer models to support policy analysis¹ is a troubled business. Studies using large computer models have met with numerous difficulties, including problems with verification, validation, sensitivity analysis, and opaqueness of model internals and outputs. These problems have often caused the justification of these models to be questioned, and the use of computer models has on various occasions been criticized as "bad science," "an excuse for not thinking," or even as "fraudulent" (Hodges, 1991; Freedman, 1981; Schrage, 1989; Stockfish, 1975). Such criticisms have been aimed at the use of computer modeling in varied domains including economics, military science, and global climatology. Even proponents of computer modeling as a methodology will admit that there are often significant problems with model quality.

The use of computer models for most policy purposes has a fundamentally different character from the paradigm of modeling in engineering and the hard sciences. In particular, it is important to distinguish between "predictive" models that may be experimentally validated to predict detailed behavior in a target system and "exploratory" models that cannot. Failure to properly recognize this distinction is the primary cause for most of the problems that have been encountered in the use of computer modeling for policy analysis. The confusion between prediction using experimentally validated models and other uses for modeling is extremely dangerous as model outcomes can be misleading when erroneously interpreted as predictions.

Enormous increases in the availability of computational power in the past few years has made the exploratory use of computer models possible for the first time. The exploratory use of nonpredictive models can usefully inform decisionmaking, but it requires a different approach to modeling than is used when prediction is possible.

The rationale of "exploratory modeling" advanced here explains how computer models can fruitfully support policy studies. It also suggests improvements both in the methodology and technology of computer modeling for policy. Where prediction and experimental validation are not possible, no single "true" model for the system of interest can be agreed upon. However, in such situations we need not be restricted to the use of any

¹Computer models have other uses, such as gaming systems in support of training or as a representation of expert knowledge. This Note does not consider these uses; it focuses solely on the use of modeling to support analytic studies.

single model in our reasoning. This Note explains how this flexibility (supported by modern computing capabilities) opens up new options for modeling complex systems, avoiding the use of large models and easing problems with verification, understandability, and sensitivity analysis.

Section 2 describes the problems to which computer models are prone and analyzes why they occur. Section 3 describes and attempts to justify the paradigm of exploratory modeling. Section 4 considers the implications of these ideas for the methodology of building models and for desirable supporting computer technology. Section 5 relates points made here to developments in computational science generally.

2. THE PROBLEM: PRETENDING TO DO WHAT CAN'T BE DONE

THE SHORTCOMINGS OF COMPUTER MODELS

Consider the following fictional account.

In 199x, the Joint Chiefs decide they need to develop improved means for making (and defending) procurement decisions. (This story would not change much if the problem was avoiding global warming or preventing the next economic recession.) In order to analyze the differential impact of alternative decisions on potential combat outcomes, they decide to build the penultimate combat simulation. The wisest experts in military science are drawn together to define the model. A crack team of programmers is assembled to implement it on the most advanced computers, using state-of-the-art software tools. All relevant databases are made available for the effort. In all cases, the best regarded modeling techniques are used, and military experts are consulted to ensure the realism of each submodel. In order that the penultimate model be valid for a wide range of contingencies, all phenomena that might potentially be influential on battle outcomes are included, and all details that might prove pivotal are represented. This results in a very detailed model (because of a nail, the shoe was lost, because of the shoe the horse, etc.). The model is seen as realistic because it includes many factors, and lots of hard engineering data can be used, enhancing the model's credibility.

The resulting computer program is, of course, quite large, involving several million lines of code. And because new studies often suggest needed modifications, the length continues to grow even after the model becomes operational. For the model to execute within a reasonable time it must run on the newest massively parallel supercomputers. Even on these machines, only a few cases can be run. The main constraint on how fast cases can be run is not execution time but the time required to set up the hundreds of thousands of input parameters. The outputs of the model are so voluminous that figuring out who is winning is not trivial either. These problems are met by yet more computer programs that automate the process of setting up initial conditions and

summarizing the outputs. The effective study turnaround time remains long, however, because the computer software requires highly trained operators, and they unfortunately are in short supply. It is still difficult to understand why certain simulated results occur, but warfare is a complex business, and no one really expected simulation to make that complexity go away. In spite of these problems, the state-of-the-art graphics makes for great demonstrations and study outputs are very compelling to their sponsors. All of the software managers and military action officers get promoted.

After some time, however, the penultimate model begins to develop enemies. Outputs often tend to show little impact for some types of forces or weapon systems, even though proponents may consider them crucial. Upon examination, certain aspects of model output can be demonstrated to be unrealistic. This is ascribed to details that failed to make it into the model. Some assumptions about the nature of warfare that were made in designing the model are contested by detractors. Although there are far too many inputs to do a thorough sensitivity analysis, occasionally someone discovers a case where a small change creates a big swing in outcome. This is pointed out to be not unrealistic, but is distressing to those who are rationalizing multi-million dollar expenditures by comparing simulation runs. In response to these problems the model is frequently revised. Unfortunately, its size make revision nontrivial, and making sure revisions have not created obscure bugs is very time consuming (unfortunately, some of the original programmers have moved on to new jobs). Eventually the entire enterprise collapses under its own weight, and use of the penultimate model is abandoned. Work immediately begins on its successor.

This story is admittedly dramatized and oriented to a worst case. Nonetheless, there exist failed simulation models whose story is quite similar, and most large simulation efforts for policy analysis have encountered at least some of the problems mentioned here. Wise analysts are of course not as foolish as the protagonists of this story. However, the effort to build the penultimate combat simulation model is completely consistent with existing conventional wisdom, evidenced both by the way we build models and the way we talk about them.

Recently, many have voiced their criticisms of big modeling efforts in policy areas (Hodges, 1991). These critics generate lists of problems to which computer models are prone. The following is yet another such list:

- Computer models often tend to be large and to continue growing throughout their history.
- It is difficult to verify that the program correctly implements the advertised conceptual model. The sheer size and complexity of models often make it essentially impossible to guarantee that they have been completely debugged and to ascertain whether there are conceptual errors in the model the program implements. These problems apply equally to the computer program and to its outputs, and affect all phases of a model's history — its construction, use, maintenance, and modification for new purposes.
- The extreme size and complexity of these models make it difficult to examine the computer code to learn what the model actually does. This problem can be somewhat mitigated by documentation, but documentation tends to be inadequate and out of date.
- Because of this opacity, experts not associated with a model must rely upon a "priesthood" of model cognoscenti. One must trust the "priesthood" to have done a good job and to be correctly portraying the details of the model's internal structure.
- Consequently, computer models are seldom subjected to peer review.
- It is extremely difficult to adequately determine how sensitive model outputs are to uncertainty in the inputs.
- There is a corresponding tendency to underestimate (or ignore) the uncertainty (or inaccuracy) of inputs to models. This includes both explicit inputs and assumptions made in the process of building the model.
- There is a strong tendency to model in detail phenomena for which good models can be constructed, and to ignore phenomena that are difficult to model, producing a systematic bias in the results.
- These technical difficulties can interact with psychological or bureaucratic tendencies to produce a host of problems, including using models to rationalize institutional prejudices, poor models driving out careful thinking, and tending to emphasize those aspects of a problem that can best be simulated. The result can

often be that models provide an illusion of analytic certainty for problems that are not that well understood, or in the worst cases provide scientific costume for points of view that are self-serving.

While these problems have varying technical attributes, and particular problems may be managed through technological improvements, the myriad complaints are but symptoms of a single fundamental problem that lies in our assumptions about what models are and how they are used.

WHY MODELS GO WRONG

The preceding list of ills has often been noted and has resulted in a variety of prescriptions, falling into two broad classes: proposed improvements to the computer technology for constructing models, and critique of the methodology of using large models. Both points of view have their merit, but both fail to capture the crux of the problem.

Technological fixes alone will not solve these problems. Computer science has provided numerous tools¹ to facilitate the construction, maintenance, modification, and verification of large programs. They are all useful and improve our ability to effectively produce large models. Further technological innovations are desirable and continue to be made. However, although the use of advanced computer technology ameliorates some of the problems listed in the previous section, it does not alter the overall pattern. A given model may be more tersely and understandably expressed in an appropriate language, but there is a threshold beyond which a program becomes difficult to understand. It is better to write computer models in an appropriate high-level language, but regardless of what language they are written in they still tend to grow large and to become unmanageable. High-level languages and other techniques will not affect the overall pattern unless the reasons the models become so large are addressed. Furthermore, no technological fix has yet been proposed for the problems of validation and sensitivity analysis. Computer technology by itself cannot resolve the general problem.

The reverse point of view is opposition to the use of computer models or criticism of instances of their use. In its most extreme form this view holds that the construction of large models is a bad idea, period. Small models have few of the ills of large ones. A small model may have limits to its utility, but a modest model whose shortcomings are easily

¹Examples are improved programming languages, improved computer/user interfaces, and managerial techniques for controlling large software engineering efforts.

established is preferable to an elaborate model with enormous potential to mislead. However, there are problems whose inherent complexity cannot be wished away, and the dictum that only small or simple problems can be addressed by computer modeling means that computers may not be used to study just those phenomena for which their potential utility seems highest. Thus, this view throws out the baby with the bath water. While there are certainly problems for which modeling is not a useful tool, the key question is not whether to build computer models, but rather how models can support policy studies.

By far the most telling evidence against either the technology or the size hypothesis for modeling problems is the striking counterexample that there are very large models implemented with relatively primitive tools (e.g., FORTRAN) that have been enormously successful. Examples include digital wind tunnels and codes developed to simulate nuclear explosions.

What differentiates those problems for which computer modeling has been strikingly successful from those for which it has been problematic? There is one clear difference: computer modeling has been successful in those domains for which the models make predictions that can be verified by experiment.² Modeling efforts become troubled when there is no possibility of experimentally validating model correctness and utility. Validation may not be possible because the necessary experiments cannot be carried out, historical data are inadequate, theory is insufficiently mature to suggest models capable of making predictions; because cases of interest require initial conditions or boundary conditions that can only be guessed at; or because nonlinearities in the model cause even modest uncertainties in the inputs to produce substantial uncertainties in the results. Most systems of interest for policy analysis cannot be predicted for several of these reasons. It is not possible to stage World War III several times in order to resolve questions of interest to combat modelers, nor to guess what the initiating circumstance might be.

Means other than predicting the outcomes of experiments have been used as validation measures, including experimental validation of submodels where possible, determining model parameters from validated sources, and the prediction of general characteristics of system behavior (as opposed to the prediction of specific details). Although these are useful checks, they result in only partial validation of the model, reducing

²By experiment I mean any use of data obtained from the target system to check the outputs of a model. Thus a model might be validated by replicating historical data. The experiment in this case is setting the model inputs to the historical situation and seeing whether the model correctly "predicts" the historical outcome.

the range of uncertainty associated with its products but still leaving it inadequate to make detailed predictions.³ Reducing the amount of uncertainty associated with a model is obviously desirable. Improving the scientific basis of models used for policy analysis is an important goal. For most such models, however, the remaining uncertainty will always be significant. As long as a model has significant uncertainty, it belongs in a separate class from models that reproducibly make detailed predictions of system behavior. Thus it is important to distinguish between experimental validation and other measures that do not (completely) validate the model. Models that are not predictive must be used much differently from those that can be (reproducibly) experimentally validated.

TRYING TO PREDICT THE UNPREDICTABLE

Building models (whether computerized or not) that make accurate predictions is much of what science is all about, and an orientation toward predictive modeling is deeply rooted in the cultural heritage of science. Consequently, in applying computer modeling to problems in policy analysis, aspects of the predictive modeling paradigm are generally employed (usually implicitly) even when the impossibility of experimental validation is well recognized.

The predictive modeling paradigm is to develop a model that captures the causal dependencies of the target system and thus yields both an improved understanding of the nature of that system and an ability to make predictions about the target system's behavior. The process of building predictive computer models has clear stages:

- Development of the theory or conceptual model
- Implementation of the corresponding computer model
- Verification that the computer program correctly implements the conceptual model (debugging)
- Validation of the model by successfully predicting target system behavior (experimental validation)

³Note that "validity" is commonly used in two different senses: (1) as a boolean (a model is either validated or it is not), and (2) as a scalar (a model may become gradually more valid as various partial checks are made). Confusion between these two uses has clouded many discussions of the validation issue. Typically the unvalidated state (in the boolean sense) of models is noted, and partial validation measures are then prescribed as an available solution.

- Sensitivity analysis to yield guidelines of how much error should be expected in the predictions given known amounts of error in the model inputs
- Using the validated model to predict where experimental determination of the answers is impossible, expensive, or otherwise unavailable (running cases).

This paradigm for modeling a target system is such a deep part of our culture that it may be regarded as "obvious." However, when experimental validation is impossible, its use is not justified. Whereas the hard sciences build models (whether implemented on computers or not) that may be verified experimentally, policymakers are typically concerned with systems for which experiments cannot be conducted, and that have characteristics making detailed prediction difficult or impossible even should experimentation be possible. Consequently, the simplistic use of modeling to predict an outcome based on its inputs is ill-suited to many problems of importance. The methodology of predictive modeling (construct a model, use its predictions to distinguish among alternative choices) is similarly ill-suited for these problems. Building models can still be of value, but to be useful the models must be used differently than they would if they made predictions. Their use requires a different epistemological rationale and motivates a different development style.

Currently, models for policy studies are built as though they were to be used to make predictions, even if that is not the fundamental intent. This assumption reveals its presence through the inclination toward single monolithic models incorporating all known factors of importance and through an orientation toward best estimate cases. This development approach does not make appropriate use of the model impossible, merely difficult and hence less likely. The (often unconscious) use of the predictive paradigm for cases where it is not appropriate causes many of the problems that have been encountered with computer modeling. Consider, for example, the prevalent tendency to emphasize "building a model" over carefully thinking through an analysis. Emphasis on building a model would make sense if the possession of a model allowed prediction. But where prediction is not possible, a model is of little use without an analytic framework that makes its outputs relevant. Particularly revealing are difficulties with verification and validation. When models can be used for prediction, a single experiment in which the model successfully predicts system behavior can provide a great deal of validation and verification. The issues of verification and validation become much more vexing when model outputs cannot be checked by experiment.

The tendency for models to grow very large is at least in part a product of the confusion between models built for prediction and those used for other purposes. When a model can be used to predict, it can be validated through its predictions. Among competing predictive models, the most preferred is the simplest that makes correct predictions. However, when models cannot be validated in such a direct fashion, the quality of a model must be assessed by other means. Often, judgments of model quality are based upon the degree of completeness (the inclusion of all factors and phenomena that might influence outcomes for at least some cases). In contrast with models of predictable target systems that are simplifications of reality, models of unpredictable systems often are attempts to copy the full complexity of the target system. The designers of such models have fallen into a worship of false-reductionism: the more details a model contains, the more accurate it will be. This reductionism is false in that no amount of detail can provide validation, only the illusion of realism. This tendency is reinforced by an economic consideration: a model is more useful if it can be applied to a wide range of cases, so extra factors must be included to extend the range of the model. But is a model that does not predict a large number of cases really more useful than a model that does not predict only a few?

The growth in model size is related to an assumption that all details must be accounted for with a single monolithic model. When an experimentally validated model can be produced, this assumption can be justified. When experimental validation is no longer possible, no model can be asserted to be a correct model of the target system, and no conclusions can be safely drawn without examining alternative models. So, where there is no single "true" model, it is no longer necessary to use one model only. As will be described below, this realization opens up new options for modeling complex systems, avoiding the use of large models and easing problems with verification, understandability, and sensitivity analysis.

Predictive models make great screwdrivers, but pound nails rather poorly. For problems involving practical barriers to experimental validation, significant uncertainties, or strong nonlinearities, a different tool is needed.

3. EXPLORATORY MODELING

THE SEARCH FOR INSIGHT

"The purpose of computing is insight, not numbers."

Richard Hamming ¹

That we may be unable to predict the behavior of a system does not mean that we know nothing about it. Similarly, a model's inability to make predictions does not necessarily make it valueless.

How then can models be used if they are not validated and cannot make reliable predictions? This question is a version of the more general one: "Of what use is partial information?". For some questions, partial information may be useless, but for many problems partial information can provide partial answers. For most policy problems, some decision must be made (at least the decision to do nothing), regardless of the level of uncertainty. Policy analysis requires understanding the implications of what is known, which for systems with significant nonlinearities may not be all that obvious. When dealing with complex systems, both that which is known and that which is uncertain may be best represented by computer models. Thus, computers can have a role in revealing the implications of what is known or believed and the possible consequences of that which is unknown or uncertain. What is at issue is not whether they can be useful but rather, how are they best employed?

To usefully employ nonpredictive models, a methodology must be adopted that is much different from that appropriate for prediction. Small numbers of cases run on a single best estimate model are informative if the model makes valid predictions, but can be very misleading otherwise. When models are used to calculate the implications of what may be assumed or hypothesized, careful study may require a multiplicity of models and cases. A predictive model is an artifact that, once it is experimentally validated, can be widely applied. By contrast, exploratory modeling serves as an infrastructure to support reasoning.

The computer can have a significant role augmenting human reasoning because of the nonlinearity of most models of interest for policy problems. For nonlinear models, it is often

¹Hamming, 1962.

difficult to anticipate how a model will behave when it is executed, and even quite simple models are capable of exhibiting behaviors that surprise their creators. Consequently, building and executing models have the potential for revealing unanticipated implications of our assumptions. Even when a model is not validated, it can serve as an "inference engine," showing us where innocuous appearing assumptions lead to predicted behaviors at variance with initial expectations. By throwing light on obvious but treacherous assumptions, computer modeling can perform important service in the search for insight. However, establishing the relevance of any such insight to particular problems must be accomplished outside of the models.²

When used for exploratory modeling, the computer functions as a prosthesis for the imagination, allowing the discovery of novel explanations of known facts or unrealized properties of conjectures. The use of this prosthesis to browse through the space of plausible models can result in improved insight. However, policy may require that this tool be used in a structured fashion.

We can pose the general problem of reasoning with incomplete information as a search problem. The goal of the search is conclusions that can be safely drawn in spite of imperfect knowledge. Such conclusions correspond to facts or relationships that are invariant across all plausible models. The search may be through the set of possible models consistent with what is known. Or it may be made at a more abstract level, where the search is among alternative formulations, with the goal of a perspective that simplifies the problem to reveal the desired invariance. Exploratory modeling uses computer models to support the search process.

During the course of an analysis, initial explorations to develop insight should lead to a structured sampling of the space of plausible models and cases. Possible approaches to such structure include a compilation of the range of possible outcomes given feasible policy options, the construction of risk/benefit tradeoffs, flexibility against the unexpected, and prudent hedges. It is important to note that unlike predictive modeling, in exploratory modeling the outputs of models have meaning only when seen in a context provided by such an analytic strategy.

²Consequently, it must be emphasized that computer models must be regarded as decision aids, not decisionmakers, in policy contexts.

SENSITIVITY ANALYSIS

It may be illuminating to consider the relation of exploratory modeling to the concept of sensitivity analysis. For any numerical computer program, sensitivity analysis is the process by which uncertainty in inputs is related to uncertainty in outputs (Ronen, 1988; Suri, 1989; Suri, 1987). For predictive models, sensitivity analysis is important because it allows possible errors in inputs to be translated to possible error ranges in outputs when running cases. For models that are believed to be predictive (because of confidence in the correctness of the conceptual model) but that lack thorough experimental validation, sensitivity analysis is critical. Extreme sensitivity of the outputs on the inputs could cause the model to be nonpredictive for practical purposes, even when the conceptual model is correct. Unfortunately, for a model of more than modest size, an exhaustive sensitivity analysis through running excursions is combinatorially impossible. Consequently, in most model building enterprises, sensitivity analysis (when performed at all) is done only as a spot check. Lip service is paid (one morally should do some sensitivity analysis), but because it is impossible to do a thorough sensitivity analysis, one does the best one can and then moves on. The models are not designed with sensitivity analysis in mind, rather they are designed as though the sensitivities were known a priori to be strongly bounded (that is, it is assumed that the model will be predictive).

For "well-behaved" models, ranges of inputs are mapped to ranges of outputs and a thorough sensitivity analysis can be made by testing extremal points. Arbitrary nonlinear models may not be so well behaved, however. A continuous range of an input value can be mapped (as a result of bifurcations or "catastrophes") to an arbitrary number of discontinuous ranges of output values. For complex models whose mathematical characteristics are not well understood, completely characterizing the implications of uncertainty in one input could require very large numbers of cases. Thus, the common tendency to perform a handful of excursions on a complex computer simulation and label that a sensitivity analysis is rather suspect.

In the context of exploratory modeling, "sensitivity analysis" is rather a misnomer. Although such an analysis can be conducted on an individual model, there can be many different models involved in the exploration of the implications of alternative assumptions and formulations. Furthermore, for exploration, sensitivity analysis is not a "nice-to-have" side task, it is the main result. Describing ranges of outcomes (assigning error ranges to the outputs) is a central goal of the analysis.

A complete sensitivity analysis assesses the behavior of all plausible models. To do this thoroughly will in general be impossible. Instead we must rely upon a strategy for sampling the (generally infinite) space of models and cases. If the statistical properties of all parameters were known, and the dependence of the outputs upon the inputs were sufficiently well behaved, a mathematically rigorous strategy for sampling could be devised. In general, however, a sampling strategy will need to be devised to search for critical cases, using human judgment to prioritize the investigation of the uncertainties involved. The result will not be a mathematically rigorous answer, but rather an imperfect image of the complete envelope that improves gradually as more cases are run.

In practice, only a finite number of cases can be run. Given a fixed analytic budget (in dollars, people, or time), analysis must provide the best insights possible based on what is known about the problem at hand. The exploration strategy will usually take advantage of knowledge not contained in the models and will be structured by the pursuit of a compelling argument illuminating the choice among policy options. Exploratory modeling allows for the flexible allocation of resources (human as well as computational) to those aspects of the problem that are judged most critical.

USING EXPLORATORY MODELING TO SUPPORT POLICY ANALYSIS

We resort to computer calculation when the complexities of detail in a reasoning process exceed our ability to do the calculation in our heads. In an earlier era, when the complexity grew too great to be kept in one's head, one relied upon a blackboard or piece of paper to hold the details. Yet no one viewed an argument as persuasive because it came from a blackboard. Arguments needed to be persuasive in and of themselves, and the blackboard was a useful tool in arriving at a persuasive argument. Similarly, I believe that the use of computers in support of policy analysis should be as "dynamic blackboards." They can be very useful in working through the details, but the final result must be compelling and inspectable on its own. Convincing via the outputs of what is ultimately a "black box" is superstitious behavior that must eventually cease as our sophistication in the use of computers to assist reasoning grows. Models that are not "black boxes" can be used for convincing, but their value is still relative to a logical context established outside of the model. The argument is the central result, not the outputs of any computer model.

Insights developed through exploratory modeling need not necessarily be private or subjective. Exploration can be completely unstructured, with the sole goal of developing intuition, or it can be guided by a strategy for surveying the envelope of plausible models in

support of a policy study. The product of such an exploration will not be a computer model, validated or otherwise. Neither will it typically be a single number, graph, or table that gives the answer to a specific question. Rather, the final product of a successful exploration will be an argument or chain of reasoning. This chain will start from assumptions, lead to some conclusion, and possibly have some model outputs as individual steps. It is quite unlikely that this argument will ever be entirely computer generated. In fact, the product of exploratory modeling could often be a line of reasoning that makes no reference to computer models whatsoever.

When models do appear in a final product, they will be providing links in the chain of reasoning: if we assume X, then Y follows computationally. The logical coherence that makes many individual steps lead to a conclusion must be provided outside of the models. Thus in constructing such an argument, models must be built and used in service to an analytic strategy, and in the conclusions of a study they are relevant only in the context of an argument that takes their limitations into account. The unaugmented outputs of a nonpredictive model never have a particular meaning for policy.

Thus, whereas experiments with computer models can produce insight, insight is required prior to the structured use of modeling, if problems arising from complexity and uncertainty are to be managed. Proper structuring of the problem space can provide a strategy for exploration designed to confront problems with uncertainty and reveal consequences for policy. Because thorough search of the plausibility envelope is not possible, such strategies will often focus on the elucidation of critical cases that support choosing one policy from a list of options even when the exact range of outcomes is unknown. Examples are the discovery of taxonomies of worst cases that allow choosing among options by risk aversion, the use of a fortiori arguments, and identifying cases where expenditures make a difference or where competing models may be distinguished. The existence proof provided by a single plausible model can demolish an erroneous "common sense" argument, and thus contribute significantly to clear thinking.

The exploratory approach compares positively with common ideas about the use of computer models in several ways. In changing from a predictive to an exploratory paradigm, we can better support the determination of expected ranges of variability of outcome, as opposed to generating best estimate predictions, which are meaningless without error estimates. By moving the focus from "the model" to the line of reasoning, numerous irrationalities in the use of models can be discouraged and opportunity is provided for the inclusion of other analytic tools and expert judgment, which is often disenfranchised by the

use of complex "black box" models. A more up-front portrayal of the exploratory aspect of the analytic process will provide for client education and reduce the likelihood of the misuse of model outputs and study conclusions.

In the past, retaining a predictive modeling orientation was partially motivated by constraints on computational experimentation imposed by limitations in computer power. However, the increasing speed of commonly available machines together with recent advances in software technology now make the aggressive exploitation of exploratory modeling possible for the first time. This paradigm can potentially provide two general advantages to policy analysis. One is a basis for improved understanding of what computer models do and do not provide, offering at least a partial remedy for existing problems with the misuse of computer modeling. The other is improvement in the utilization of computational resources in the service of policy studies. The range of uses for exploratory modeling is limited not by problems with validation, but with limitations in the flexibility of existing models and modeling support environments. Possibilities for improvements are considered in the following section.

4. IMPLICATIONS FOR BUILDING EXPLORATORY MODELS

Whereas exploratory analysis can be done using the single monolithic models that typically result from model development in the predictive style, many aspects of these models make exploratory use difficult. Furthermore, many shortcomings in the use of models in policy studies result from a lack of caution regarding the predictive capabilities of computer models. This suggests that approaches to model development that abandon the pretense of predictive modeling and provide computer support for exploratory use could provide enhanced capabilities. This section considers possible innovations in the methodology and technology of model building that could provide for greatly improved capabilities for exploratory modeling. Proposing new methodology and technology involves some speculation, but none of what is proposed below is unprecedented, and the potential benefits could be significant.

QUESTION-DRIVEN MODELING VERSUS DATA-DRIVEN MODELING

The predictive paradigm for model development is driven by the inputs to the model. Models are defined by the entities in the target system that are to be represented, their relationships, and (sub)models of their state changes over time. For predictive purposes, the primary constraint is the availability of data. Once a model is experimentally validated, its (correct) predictions can be used to address a variety of questions.

In contrast, exploratory modeling is best served by question-driven model design. For problems in which no amount of data will suffice, the place to start is not with how to represent aspects of the target system (what data the model contains) but rather with how the model will be used (what answer the model is to provide). To use computer models to inform policy, the strategy for how they will be useful should be factored into the design. The questions that need to be answered can have a profound impact on what representations are desirable. The realization that the way the models are actually used is exploratory rather than predictive (at least if they are used responsibly) opens up the option of designing models top-down (driven by their use) rather than bottom up (driven by the entire set of data items that might be relevant to some question). Recent progress in software engineering allows for the possibility of efficiently pursuing a much more interactive modeling style.

At present, most model development betrays an assumption of predictive use in its sequential "two-phase" style. The first phase is to construct a general model of the target system. The second phase is to use that model to answer various questions by performing runs. This approach makes sense if the model can be validated and used for prediction; however, it creates enormous barriers to facile exploration. Exploratory use will require a variety of modeling assumptions to address various needs. The "two-phase" style produces single models that may not be well designed for the purposes they are put to. Such compartmentalized development tends to result in large models (a result of trying to anticipate all possible uses) that present serious barriers to adaptation for new uses.

Any single model makes some questions relatively easy to ask and others essentially impossible because of the structure of the model. As the model grows in size, all questions become proportionately harder to ask. Models built prior to identifying what the interesting questions are will seldom be ideal (or even adequate) to address those questions. The use of a large multipurpose model thus carries with it the enormous risk that the model will constrain the set of questions that are asked, creating a systematic bias to the analysis. This is akin to looking for a lost quarter only where the light is good.

NOT ONE MODEL, BUT MANY

Since many key questions may be identified during the process of an analysis, exploratory analysis is best pursued using models designed and constructed in the context of the analytic task and revised iteratively during the course of the analysis. This requirement creates a need for different software tools, and also raises issues regarding the methodology of model development. Rather than developing a single model that will be validated by predicting the outcome of experiments, exploratory modeling must be iterative¹ (models are frequently redesigned and reimplemented) and adaptive (models are revised as a result of what is learned through their use).

Since no model generated will be a valid representation of the true state of the world, there is no need to produce a single model. Consequently, the complexities of the target system could be represented by an ensemble of models and by supporting analysis. The scope of any individual model can be designed to maximize its utility for answering a specific question (including the need for understandability and sensitivity analysis). The paradigm of exploration allows for more flexible adjustment of model structure and content

¹See Boehm (1988) for a discussion of iterative modeling in a broader context.

to serve the goals of a particular analysis. Better studies result if modeling is made to serve study goals, rather than studies being driven by what can be modeled. The linear process of developing a conceptual model, implementing it as a program, and running cases must be replaced by a much more iterative approach, where preliminary modeling leads to insights that affect the design of subsequent models.

During the process of exploration, many questions may be asked. Taken to extremes, this approach could result in a separate model for each question. All else being equal, a smaller number of versatile models are to be preferred as they will require less effort in their construction. Thus we must seek a tradeoff between ease of construction (implying few models) and ease of appropriate use (implying many). While this tradeoff will depend upon the specifics of the problem (as well as model design), the various problems attending the use of large models suggest that current methodology may be far from optimal.

The capability to produce numerous models can be provided in a number of ways. Parameterized models allow for the exploration of the impact of ranges of parameter settings (and thus provide multiple logical models with a single piece of software). Structural model variations can be explored through model revision. Improving software tools for rapid prototyping can be helpful in making these revisions possible. Instead of building megamodels to support a variety of studies, modeling environments could be constructed that incorporate baseline models, libraries of model components, and other tools to aid in model construction, so that the process of building numerous model variants is made tractable.

SELECTIVE RESOLUTION

Exploratory modeling allows for much greater flexibility in choosing appropriate levels of resolution in models. The model that is built to answer a particular question should generally be the smallest (lowest resolution) that satisfies that purpose. Keeping the model as limited as possible minimizes problems with understandability and sensitivity analysis. As different questions are asked during the course of an analysis, models of different resolution may be required. Addressing broad tradeoffs may require aggregated models of wide scope, whereas models for specific questions may require more focus and detail.

The resolution of a model must be distinguished from its size. Commonly, higher resolution models are larger, as they attempt to address the same range of question while providing greater detail. When this approach is taken, however, model understandability and sensitivity analysis must be sacrificed. For exploratory use, we may entertain keeping

model size within bounds (maintaining understandability and sensitivity analysis) by trading off resolution for breadth. Typically, broader questions will be posed early in an analysis, resulting in highly aggregated models. The insight gained from working with broad aggregated models can motivate specific questions, requiring detailed modeling guided by results obtained in preliminary analysis.

The selective refinement of aggregated models into higher resolution ones provides an option for a structured approach to the problem of sensitivity analysis and the exploration of the envelope of plausible models. Initial aggregated modeling efforts can be used to bound and give insight into the question of range of plausible outcomes. The limited size of aggregated models facilitates extensive sensitivity analysis. Subsequent modeling incorporating more details would be guided by the outcomes of prior sensitivity analysis. Efforts could be focused on adding those details or building the additional models that contribute the most to the quality of the answer to the question at hand. Low-resolution models (which typically run quickly) can be used to discover critical cases to be run at higher resolution (which can be more computationally intensive). The results of high-resolution modeling may also suggest parameter choices for more aggregated parameterized models.

It will generally be desirable to increase the resolution of a model only for those parameters that are shown to be critical to the question at hand. In this way, the results of preliminary analysis with aggregated models can guide the allocation of resources in more detailed modeling. At the same time, by adding resolution only where necessary in the context of a specific question, the use of monolithic high-resolution models, with all their attendant difficulties, may be avoided.

COMPUTER TECHNOLOGY TO SUPPORT EXPLORATORY MODELING

Harnessing computational power in service to human creativity requires that the computer support be carefully designed to supplement and not inhibit. Addressing complex and uncertain questions through exploratory modeling requires appropriate support software.

It is interesting to note that the "two-phase" approach to model development is reminiscent of the rigid style of program development characteristic of the batch-processing era of computation. Modern developments in interactive software environments allow for a more fluid approach to computation, which is characterized by faster development times and greater human productivity. This suggests that computer modeling for policy analysis is not receiving the full benefit of appropriate computer support, and that a software environment

designed to support exploratory modeling could provide significant improvements over current approaches.

The methodology proposed here would in many cases replace a complex monolithic model with numerous simpler models. Relative model simplicity (and hence many other desirable qualities) would be obtained by exporting complexity outside the model and into the surrounding computational environment and analytic context. One reason why this shift is attractive is that when complexity is not buried inside a procedural model, more powerful software tools can be provided to manage it. Aggressive exploitation of the approach championed here requires the development of such tools. Whereas the usefulness of tools must be demonstrated by their application to real problems, computer technology exists which if applied to the problems of exploratory modeling promises significant enhancement in computational support compared with existing approaches to computer modeling.

A computer environment to support exploratory modeling should allow smart people to efficiently navigate the space of plausible models and model outcomes to construct lines of reasoning and to make themselves smarter. Computer tools are required that meet the following three general needs:

- Support for iterative and adaptive modeling
- Assistance in managing the complexity of numerous models, cases, and relationships between them
- Means for portraying the results of exploratory modeling.

Support for Iterative and Adaptive Modeling

The predictive paradigm of model development often results in models that are effectively "black boxes". Exploratory modeling requires just the reverse: WYSIWYG² models that can be easily understood and revised. The goal of WYSIWYG modeling is made possible by the flexibility of exploratory modeling to obey constraints on the size of models, the availability of high-level programming languages designed for understandability and modifiability, and the use of interactive computer software environments allowing easy inspection and manipulation of model source code, parameters, and outputs. A fully developed set of tools for easing program modification and understanding could be characterized as a computer-assisted software engineering (CASE) environment for models.

²WYSIWYG = What you See Is What You Get. In other words, everything that is important in the model is there to be seen. Models must be inspectable by their users without prohibitive amounts of effort.

The style of software development often referred to as "rapid prototyping" is well suited for the needs of exploratory modeling. A key element here is the use of high-level languages that support WYSIWYG modeling. The success of the RAND-ABEL Table features (Allen and Wilson, 1988; Shapiro et al., 1988; Shapiro et al., 1985) demonstrates that it is possible to write models that are inspectable, understandable, and modestly modifiable by nonprogrammers. Further progress in this direction, along with other initiatives in language design, could lead to fully general modeling languages providing the accessibility of spreadsheets.

New models need not always be built from scratch; exploratory modeling may also be supported through the use of parameterized models and by combining and revising existing model components. It may be possible to construct modeling environments for specific policy areas incorporating baseline model components to allow model construction through combining model components, varying parameters, and model revision.

Constructing new models by combining model components requires standards for model interfaces. Interface standards would also facilitate construction of hierarchic ensembles of models and standard tools for viewing the behavior of models. The definition of such interfaces is a challenging problem requiring innovation. The problem is eased somewhat by having model components interact through a common data facility with an associated standard data dictionary.

A variety of hypothetical software tools could also help to simplify the work of exploratory modeling, although experimentation would be required to determine their actual utility. One interesting possibility is a tool to support the "bottom-up" calculation of the sensitivities associated with a particular model run through the automatic generation of excursions. While this would in general involve significant computation, the rapid growth of available computing power implies that a useful number of excursions could be computed for models of moderate size. Furthermore, the possibility exists that significant computation could be saved via software that caches intermediate sensitivity calculations in the context of the large number of runs required to perform the complete sensitivity analysis (Rothenberg et al., 1990).

Support for Managing the Complexity of an Evolving Analysis

Although the use of multiple models can allow models to become simpler, the relationships among models and the process of model and case management will become more complex. The overall amount of detail may be as great or greater than that in a single monolithic model, but would be distributed across multiple computer programs. In a sense, the use of multiple models enforces a strict form of modularity, with implied benefits for understandability and verification. This has the virtue of moving complexity out of computer program internals (where it can be difficult to understand) and into the declarative realm of model interrelationships where it may be more easily viewed, understood, and manipulated.³

There are various sorts of information that could usefully be represented in the computer, including:

- The evolving logical argument
- The plan for the analysis
- The various cases that support points in the logical argument
- The models that were used to analyze these cases
- The databases used in running the models
- Other sources of information (such as off-line analyses, compilations of expert opinion, historical data, or supporting graphics).

The ability of the analyst to keep track of the myriad details of model characteristics, interrelationships, cases, histories, implications, status, outcomes, and the like could be greatly enhanced by an appropriate software environment. Even where a single model is used for exploration, many cases must be run, creating the need for tools to help manage the resulting complexity. Whereas the development and running of "the model" can be the focus of activity for predictive modeling, exploration will not typically be model centered, but rather is driven by the evolving analysis. A computational environment to support exploratory modeling would similarly not be centered on a model but rather would have a representation of the evolving logical argument and plan of analysis. These would form the conceptual center for a data facility that would serve as an electronic record of the evolving

³Note that such benefits are contingent on insightful structuring of the problem. Whereas declarative representations may be generally easier to understand than procedural representations, no technique is a panacea.

chain of reasoning that constitutes the analysis. Such a facility would be more than a database, as it would contain not only data but also computer models, model runs, model outputs, human notations, and all needed relationships among these entities. Such a facility would in fact be a hypermedium (Barrett, 1988; CACM, 1988; Wurman, 1989) for modeling.

Portraying the Results of Exploratory Modeling

Individual model runs can produce voluminous data; multiple such runs can produce astronomical quantities. Developing intuition based upon the results of modeling requires adequate means for viewing these results. The deluge of data that can be generated makes it impossible for users to quantitatively examine more than a fraction of it. If insight is to be generated from these outputs, means must be available to easily view the data for various purposes. Means for viewing the results of exploratory modeling would be useful both for presenting final results to the consumer of the analysis and for providing a powerful means for the analyst to improve his intuition.

With the advent of raster graphics, entire fields of variables can be converted to color images. Information conveyed in this way undergoes a qualitative change because it utilizes the tremendous pattern recognition capabilities of the human eye-brain system. An environment for exploratory modeling should include capabilities for the visualization of data harvested across multiple cases.

Useful facilities may include not only graphical displays but also textual data presentation tools and statistical facilities for the summarization of the results of multiple cases. For example, regression analysis could be used to generate a simplified model capturing the variability observed in multiple experiments with more complicated models (e.g., repro models (Goeller et al, 1985)). Where the range of plausible outcomes is too complex to be expressed simply by a number or a graph, but involves sets of tradeoffs the decisionmaker should be sensitive to, a spreadsheet model synopsizing the relationships could be a useful deliverable. A modeling environment could augment traditional forms of conveying results with simple computer models. The ideal environment would allow the study results, including final reports, briefing slides, and deliverable models, to be generated in the same environment as the analysis. Once an analysis was complete, that part intended for the end user could be detached from the supporting material and transferred via floppy disk or CD-ROM.

An Environment for Exploratory Modeling

The ideal computer environment for exploratory modeling would constitute a "dynamic blackboard," allowing users with problem-smarts but only moderate computer expertise to explore a universe of alternative problem formulations and computed implications. The execution of models is only one aspect of this environment—recording and helping to organize the growing body of results will be vital. Such a history would be a record of all modeling experiments, including the model variants, and data going into any modeling experiment as well as its outcome. The ultimate environment would have attributes of brainstorming and outlining tools, database facilities, version control systems, and general-purpose modeling environments. Such an environment would assist the user in keeping track of an evolving analysis involving the construction of multiple models and model variants, case runs, changing assumptions used, and tentative conclusions drawn.

The core representation tying all aspects of the environment together would be the "story" of the analysis. This story would at the beginning of the analysis be the plan for the analysis; at the end it would be the completed argument of assumptions leading to conclusions. During the course of the analysis, the "story" would be a mixture of the two, and would generally be considerably revised during the course of an analysis. Particular cases and models to support them would be justified by steps in the "story." Thus the story would not only be a human readable representation of the (notional or incomplete) study results, it would also serve as the root of a complex network of interrelated entities, including (descriptive) conceptual models, (executable) computer models, input data, case runs, model outputs, checkpoint model states, and assorted others.

Although no computer software can guarantee good work through its use, a support environment providing assistance for those tasks that computers do well (i.e., keeping track of details) could free talented users to do what they are best at, discovering meaningful order among a myriad of details.

5. SCIENCE, OPERATIONS RESEARCH, AND THE ROLE OF THE COMPUTER

The computer is a powerful tool, but it can mislead as well as illuminate. Initial attitudes toward computers have often manifested a naive belief in computer superiority, free from human error. With experience this is replaced by an appreciation that a computerized model allows one to view the implications of the model, but does not provide any special access to truth. In addition to uncertainties regarding input data and the correctness of the conceptual model, there is the additional uncertainty of whether the machine (software and hardware) is performing correctly. Rather than a source of specially accurate results, the outputs of computerized models must be viewed with additional suspicion.

However, the computer's capabilities for rapidly performing many more arithmetic or logical operations than the human mind gives it a prominent role in addressing problems of great complexity. We are still early in the process of understanding how best to design computer systems so that human capabilities are enhanced, not eclipsed, and the strengths of the computer utilized and its liabilities minimized.

The initial use of computers in many scientific fields has been primarily for data reduction and predictive modeling. The increasing availability of computational power has resulted in the adoption of exploratory modeling approaches by researchers in many fields (Anderson, 1988; Campbell et al., 1985; Lipman et al., 1989; Rose and Dobson, 1985; Strauss, 1974). Exploratory use involves the "guessing" of details of systems for which there are no data (such as the behavior of subatomic particles at very high energies or the spatio-temporal activation patterns of large numbers of neurons in the brain). The implications of these guesses can be computed, allowing the computer-assisted researcher to look for interesting "guesses." As computing becomes easier than performing experiments, this style of use becomes increasingly attractive. Its introduction is typically accompanied by controversy, but exploratory approaches have been gaining credence as increasing numbers of workers in various fields become computer adept. An example of scientific discovery based upon exploratory modeling is the development of chaos theory, in which anomalous behaviors were first seen in computer simulations, which led to later studies of physical systems and mathematic theory. The use of exploratory modeling to "break trail" for more traditional science is likely to become increasingly important. The requirement that a model be validated prior to its use was motivated by computationally impoverished

conditions. Now that reverse conditions hold, the use of the computer as a prosthesis for the imagination is increasingly viable.

No discipline seems better disposed to benefit from the exploratory approach to computation than the policy sciences, due to the complexity of the systems being reasoned about and the abundance of problems for which no model can be experimentally validated. No technical innovation is likely to be a panacea; however, the intelligent delivery of computer power to support policy analysis could contribute significantly to improved decisionmaking for complex and uncertain problems.

BIBLIOGRAPHY

- Allen, Patrick D., and Barry Wilson, "Modeling Qualitative Issues in Military Simulations with the RAND-ABEL Language," *Proceedings of the 1988 Winter Simulation Conference*, 1988.
- Anderson, Alan, "Learning from a Computer Cat," *Nature*, Vol. 331, February 1988, pp. 657-658.
- Barrett, E., *Text, Context, and Hypertext, Writing with and for Computers*, The MIT Press, Cambridge, MA, 1988.
- Boehm, Barry W., "A Spiral Model of Software Development and Enhancement," *IEEE Computer*, May 1988, pp. 61-72.
- Brand, S., *The Media Lab*, Penguin Books, New York, 1987.
- Campbell, David, Jim Crutchfield, Doynne Farmer, and Erica Jen, "Experimental Mathematics: The Role of Computation in Nonlinear Science," *Communications of the ACM*, Vol. 28, No. 4, 1985, pp. 374-384.
- CACM (Communications of the ACM), *Special Issue on Hypertext*, Vol. 31, No. 7, June 1988.
- Freedman, D.A., "Some pitfalls in large econometric models: A case study," *Journal of Business*, Vol. 54, 1981, pp. 477-500.
- Goeller, Bruce F., and The PAWN Team, "Planning the Netherlands' Water Resources," *Interfaces*, Vol. 15, No. 1, January-February, 1985, pp. 3-33.
- Hamming, Richard, *Numerical Methods for Scientists and Engineers*, McGraw-Hill, New York, 1962.
- Hodges, James, "Six (or so) Things You Can Do With a Bad Model," *Operations Research*, Vol. 39, No. 3, May-June 1991, pp. 355-365.
- Lipton, Richard J., Thomas G. Marr, and J. Douglas Welsh, "Computational Approaches to Discovering Semantics in Molecular Biology," *Proceedings of the IEEE*, Vol. 77, No. 7, July 1989, pp. 1056-1060.
- Ronen, Yigal (ed.), *Uncertainty Analysis*, CRC Press, Boca Raton, Florida, 1988.
- Rose, David, and Vernon Dobson, "Methodological Solutions for Neuroscience," in D. Rose and V. Dobson (eds.), *Models of the Visual Cortex*, John Wiley & Sons, New York, 1985.

- Rothenberg, J., N. Z. Shapiro, and C. Hefley, "A 'Propagative' Approach to Sensitivity Analysis," *Proceedings of the AI, Simulation and Planning in High Autonomy Systems Conference*, Tucson, Arizona, March 26-27, B. Zeigler and J. Rozenblit (eds.), IEEE Computer Society Press, Los Alamitos, California, 1990, pp. 10-16.
- Michael Schrage, "Why Beauty of Scientific Models Is Often Only Skin Deep," *Los Angeles Times*, 21 December 1989, p. D1.
- Shapiro, Norman Z., H. Edward Hall, Robert H. Anderson, and Mark LaCasse, *The RAND-ABEL Programming Language, History, Rationale, and Design*, RAND, R-3274-NA, August 1985.
- Shapiro, Norman Z., H. Edward Hall, Robert H. Anderson, Mark LaCasse, Marrietta S. Gillogly, and Robert Weissler, *The RAND-ABEL Programming Language: Reference Manual*, RAND, N-2367-1-NA, December 1988.
- Stockfish, J. A., *Models, Data, and War: A Critique of the Study of Conventional Forces*, RAND, R-1526-PR, 1975.
- Strauss, C. M., "Computer-encouraged Serendipity in Pure Mathematics," *Proceedings of the IEEE*, 62, 1974.
- Suri, R., "Infinitesimal Perturbation Analysis for General Discrete Event Systems," *JACM*, Vol. 34, No. 3, July 1987, pp. 686-717.
- Suri, R., "Perturbation Analysis: The State of the Art and Research Issues Explained via the GI/G/1 Queue," *Proceedings of the IEEE*, Vol. 77, No. 1, January 1989.
- Wurman, R., *Information Anxiety*, Doubleday, New York, 1989.